

WHAT IS CLAIMED IS:

1 1. A computer-implemented method of managing a storage environment
2 comprising storage units, the method comprising:
3 detecting a condition indicating that capacity utilization balancing is to be
4 performed for a plurality of storage units;
5 identifying a first storage unit from the plurality of storage units from which
6 data is to be moved;
7 identifying a file stored on the first storage unit to be moved;
8 identifying a storage unit from the plurality of storage units for storing the file;
9 moving the file from the first storage unit to the storage unit identified for
10 storing the file; and
11 repeating, the identifying a file stored on the first storage unit to be moved, the
12 identifying a storage unit from the plurality of storage units for storing the file, and the
13 moving the file from the first storage unit to the storage unit identified for storing the file,
14 until the condition is determined to be resolved.

1 2. The method of claim 1 wherein:
2 detecting the condition comprises detecting a condition that indicates that used
3 storage capacity for at least one storage unit from the plurality of storage units has exceeded a
4 first threshold value; and
5 the condition is determined to be resolved when the used storage capacity of
6 the at least one storage capacity for the storage unit falls below the first threshold value.

1 3. The method of claim 2 wherein identifying the first storage unit
2 comprises:
3 identifying the at least one storage unit whose used storage capacity has
4 exceeded the first threshold value as the first storage unit.

1 4. The method of claim 1 wherein:
2 detecting the condition comprises detecting that a difference in used capacity
3 of a least full storage unit and the most full storage unit in the plurality of storage units has
4 exceeded a second threshold value; and
5 the condition is determined to be resolved when the difference is within the
6 second threshold value.

1 5. The method of claim 4 wherein identifying the first storage unit
2 comprises:
3 identifying the most full storage unit as the first storage unit.

1 6. The method of claim 1 further comprising:
2 determining a storage unit from the plurality of storage units that is least full;
3 determining a storage unit from the plurality of storage units that is most full;
4 determining a difference in used capacity between the least full storage unit
5 and the most full storage unit; and
6 performing, the identifying a first storage unit from the plurality of storage
7 units from which data is to be moved, the identifying a file stored on the first storage unit to
8 be moved, the identifying a storage unit from the plurality of storage units for storing the file,
9 the moving the file from the first storage unit to the storage unit identified for storing the file,
10 and the repeating, only if the difference exceeds a pre-configured threshold value.

1 7. The method of claim 1 wherein identifying a file stored on the first
2 storage unit to be moved comprises:
3 generating a score for each file included in a plurality of files stored on the
4 first storage unit; and
5 selecting a file, from the files stored on the first storage unit, a file with the
6 highest score as the file to be moved.

1 8. The method of claim 1 wherein identifying a storage unit from the
2 plurality of storage units for storing the file comprises:
3 generating a score for the storage units in the plurality of storage units; and
4 selecting a storage unit from the plurality of storage units with the highest
5 score as the storage unit for storing the file.

1 9. The method of claim 1 wherein repeating comprises:
2 determining a storage unit from the plurality of storage units that is least full;
3 determining a storage unit from the plurality of storage units that is most full;
4 determining a difference in used capacity between the least full storage unit
5 and the most full storage unit; and
6 repeating, the identifying a file stored on the first storage unit to be moved, the
7 identifying a storage unit from the plurality of storage units for storing the file, and the

8 moving the file from the first storage unit to the storage unit identified for storing the file,
9 only if the difference exceeds a pre-configured threshold value.

1 10. The method of claim 1 wherein the plurality of storage units comprises
2 at least one storage unit assigned to a first server and at least another storage unit assigned to
3 a second server distinct from the first server.

1 11. The method of claim 1 wherein an original file stored on the first
2 storage unit is not moved until all migrated files stored on the first storage unit have been
3 moved.

1 12. In a storage environment comprising a plurality of storage units
2 assigned to one or more servers, a computer-implemented method of performing capacity
3 utilization balancing, the method comprising:

4 monitoring a first group of storage units from the plurality of storage units;

5 receiving a first signal indicative of a condition;

6 responsive to the first signal, determining a first storage unit from the first
7 group of storage units from which data is to be moved; and

8 moving data from the first storage unit to one or more other storage units in
9 the first group of storage units until the condition is resolved.

1 13. The method of claim 12 wherein:

2 the first signal indicates that used storage capacity for a storage unit from the
3 first group of storage units has exceeded a first threshold; and

4 determining the first storage unit comprises identifying the storage unit whose
5 used storage capacity has exceeded the first threshold as the first storage unit.

1 14. The method of claim 12 wherein moving data from the first storage
2 unit to one or more other storage units in the first group of storage units comprises:

3 identifying a file stored on the first storage unit to be moved;

4 identifying a storage unit from the first group of storage units for storing the
5 file;

6 moving the file from the first storage unit to the storage unit identified for
7 storing the file; and

8 repeating, the identifying a file, identifying a storage unit, and the moving the
9 file, until the condition is determined to be resolved.

1 15. The method of claim 14 wherein identifying a file stored on the first
2 storage unit to be moved comprises:
3 generating a score for each file included in a plurality of files stored on the
4 first storage unit; and
5 selecting a file, from the files stored on the first storage unit, with the highest
6 score as the file to be moved.

1 16. The method of claim 14 wherein identifying a storage unit from the
2 first group of storage units for storing the file comprises:
3 generating a score for the storage units in the first group of storage units; and
4 selecting a storage unit from the first group of storage units with the highest
5 score as the storage unit for storing the file.

1 17. The method of claim 12 wherein moving data from the first storage
2 unit to one or more other storage units in the first group of storage units comprises:
3 moving a first file stored on the first storage unit to a first target storage unit
4 included in the first group of storage units; and
5 moving a second file stored on the first storage unit to a second target storage
6 unit included in the first group of storage units, wherein the second target storage unit is
7 distinct from the first target storage unit.

1 18. The method of claim 12 further comprising:
2 determining a storage unit from the first group of storage units that is least
3 full;
4 determining a storage unit from the first group of storage units that is most
5 full;
6 determining a difference in used capacity between the least full storage unit
7 and the most full storage unit; and
8 performing the determining the first storage unit step and the moving step only
9 if the difference exceeds a pre-configured threshold value.

1 19. The method of claim 12 further comprising:
2 receiving information indicative of storage units from the plurality of storage
3 units to be included in the first group of storage units.

1 20. The method of claim 12 wherein the first group of storage units
2 comprises at least one storage unit assigned to a first server and at least another storage unit
3 assigned to a second server distinct from the first server.

1 21. The method of claim 12 wherein original data stored on the first
2 storage unit is not moved until all migrated data stored on the first storage unit has been
3 moved.

1 22. A computer program product stored on a computer-readable medium
2 for balancing capacity utilization in a storage environment comprising storage units, the
3 computer program product comprising instructions for:

4 detecting a condition indicating that capacity utilization balancing is to be
5 performed for a plurality of storage units;

6 identifying a first storage unit from the plurality of storage units from which
7 data is to be moved;

8 identifying a file stored on the first storage unit to be moved;

9 identifying a storage unit from the plurality of storage units for storing the file;

10 moving the file from the first storage unit to the storage unit identified for
11 storing the file; and

12 repeating, the identifying a file stored on the first storage unit to be moved, the
13 identifying a storage unit from the plurality of storage units for storing the file, and the
14 moving the file from the first storage unit to the storage unit identified for storing the file,
15 until the condition is determined to be resolved.

1 23. The computer program product of claim 22 wherein:

2 the instructions for detecting the condition comprise instructions for detecting
3 a condition that indicates that used storage capacity for at least one storage unit from the
4 plurality of storage units has exceeded a first threshold value; and

5 the condition is determined to be resolved when the used storage capacity of
6 the at least one storage capacity for the storage unit falls below the first threshold value.

1 24. The computer program product of claim 23 wherein the instructions for
2 identifying the first storage unit comprise:

3 instructions for identifying the at least one storage unit whose used storage
4 capacity has exceeded the first threshold value as the first storage unit.

1 25. The computer program product of claim 22 wherein:
2 the instructions for detecting the condition comprise instructions for detecting
3 that a difference in used capacity of a least full storage unit and the most full storage unit in
4 the plurality of storage units has exceeded a second threshold value; and
5 the condition is determined to be resolved when the difference is within the
6 second threshold value.

1 26. The computer program product of claim 25 wherein the instructions for
2 identifying the first storage unit comprise instructions for identifying the most full storage
3 unit as the first storage unit.

1 27. The computer program product of claim 22 further comprising
2 instructions for:
3 determining a storage unit from the plurality of storage units that is least full;
4 determining a storage unit from the plurality of storage units that is most full;
5 determining a difference in used capacity between the least full storage unit
6 and the most full storage unit; and
7 performing, the identifying a first storage unit from the plurality of storage
8 units from which data is to be moved, the identifying a file stored on the first storage unit to
9 be moved, the identifying a storage unit from the plurality of storage units for storing the file,
10 the moving the file from the first storage unit to the storage unit identified for storing the file,
11 and the repeating, only if the difference exceeds a pre-configured threshold value.

1 28. The computer program product of claim 22 wherein the instructions for
2 identifying a file stored on the first storage unit to be moved comprise instructions for:
3 generating a score for each file included in a plurality of files stored on the
4 first storage unit; and
5 selecting a file, from the files stored on the first storage unit, a file with the
6 highest score as the file to be moved.

1 29. The computer program product of claim 22 wherein the instructions for
2 identifying a storage unit from the plurality of storage units for storing the file comprise
3 instructions for:
4 generating a score for the storage units in the plurality of storage units; and

5 selecting a storage unit from the plurality of storage units with the highest
6 score as the storage unit for storing the file.

1 30. The computer program product of claim 22 wherein the instructions for
2 repeating comprise:

3 determining a storage unit from the plurality of storage units that is least full;
4 determining a storage unit from the plurality of storage units that is most full;
5 determining a difference in used capacity between the least full storage unit
6 and the most full storage unit; and

7 repeating, the identifying a file stored on the first storage unit to be moved, the
8 identifying a storage unit from the plurality of storage units for storing the file, and the
9 moving the file from the first storage unit to the storage unit identified for storing the file,
10 only if the difference exceeds a pre-configured threshold value.

1 31. The computer program product of claim 22 wherein the plurality of
2 storage units comprises at least one storage unit assigned to a first server and at least another
3 storage unit assigned to a second server distinct from the first server.

1 32. The computer program product of claim 22 wherein an original file
2 stored on the first storage unit is not moved until all migrated files stored on the first storage
3 unit have been moved.

1 33. A computer program product stored on a computer-readable medium
2 comprising code for performing capacity utilization balancing in a storage environment
3 comprising a plurality of storage units assigned to one or more servers, the computer program
4 product comprising code for:

5 monitoring a first group of storage units from the plurality of storage units;
6 receiving a first signal indicative of a condition;
7 determining, responsive to the first signal, a first storage unit from the first
8 group of storage units from which data is to be moved; and

9 moving data from the first storage unit to one or more other storage units in
10 the first group of storage units until the condition is resolved.

1 34. The computer program product of claim 33 wherein:
2 the first signal indicates that used storage capacity for a storage unit from the
3 first group of storage units has exceeded a first threshold; and

4 the code for determining the first storage unit comprises code for identifying
5 the storage unit whose used storage capacity has exceeded the first threshold as the first
6 storage unit.

1 35. The computer program product of claim 33 wherein the code for
2 moving data from the first storage unit to one or more other storage units in the first group of
3 storage units comprises code for:
4 identifying a file stored on the first storage unit to be moved;
5 identifying a storage unit from the first group of storage units for storing the
6 file;
7 moving the file from the first storage unit to the storage unit identified for
8 storing the file; and
9 repeating, the identifying a file, identifying a storage unit, and the moving the
10 file, until the condition is determined to be resolved.

1 36. The computer program product of claim 35 wherein the code for
2 identifying a file stored on the first storage unit to be moved comprises:
3 code for generating a score for each file included in a plurality of files stored
4 on the first storage unit; and
5 code for selecting a file, from the files stored on the first storage unit, with the
6 highest score as the file to be moved.

1 37. The computer program product of claim 35 wherein the code for
2 identifying a storage unit from the first group of storage units for storing the file comprises:
3 code for generating a score for the storage units in the first group of storage
4 units; and
5 code for selecting a storage unit from the first group of storage units with the
6 highest score as the storage unit for storing the file.

1 38. The computer program product of claim 33 wherein the code for
2 moving data from the first storage unit to one or more other storage units in the first group of
3 storage units comprises:
4 code for moving a first file stored on the first storage unit to a first target
5 storage unit included in the first group of storage units; and

6 code for moving a second file stored on the first storage unit to a second target
7 storage unit included in the first group of storage units, wherein the second target storage unit
8 is distinct from the first target storage unit.

1 39. The computer program product of claim 33 further comprising:
2 code for determining a storage unit from the first group of storage units that is
3 least full;
4 code for determining a storage unit from the first group of storage units that is
5 most full;
6 code for determining a difference in used capacity between the least full
7 storage unit and the most full storage unit; and
8 code for performing the determining the first storage unit step and the moving
9 step only if the difference exceeds a pre-configured threshold value.

1 40. The computer program product of claim 33 further comprising:
2 code for receiving information indicative of storage units from the plurality of
3 storage units to be included in the first group of storage units.

1 41. The computer program product of claim 33 wherein the first group of
2 storage units comprises at least one storage unit assigned to a first server and at least another
3 storage unit assigned to a second server distinct from the first server.

1 42. The computer program product of claim 33 wherein original data
2 stored on the first storage unit is not moved until all migrated data stored on the first storage
3 unit has been moved.

1 43. In a storage environment comprising storage units, a system
2 comprising:
3 at least one processor; and
4 a memory operatively coupled to the processor, the memory storing program
5 instructions that when executed by the processor, cause the processor to:
6 detect a condition indicating that capacity utilization balancing is to be
7 performed for a plurality of storage units,
8 identify a first storage unit from the plurality of storage units from
9 which data is to be moved,
10 identify a file stored on the first storage unit to be moved,

11 identify a storage unit from the plurality of storage units for storing the
12 file,
13 move the file from the first storage unit to the storage unit identified
14 for storing the file, and
15 repeat, the identification of a file stored on the first storage unit to be
16 moved, the identification of a storage unit from the plurality of storage units for storing the
17 file, and the move of the file from the first storage unit to the storage unit identified for
18 storing the file, until the condition is determined to be resolved.

1 44. The system of claim 43 wherein the program instructions when
2 executed by the processor, cause the processor to detect a condition that indicates that used
3 storage capacity for at least one storage unit from the plurality of storage units has exceeded a
4 first threshold value, and the condition is determined to be resolved when the used storage
5 capacity of the at least one storage capacity for the storage unit falls below the first threshold
6 value.

1 45. The system of claim 44 wherein the program instructions when
2 executed by the processor, cause the processor to identify the at least one storage unit whose
3 used storage capacity has exceeded the first threshold value as the first storage unit.

1 46. The system of claim 43 wherein the program instructions when
2 executed by the processor, cause the processor to detect the condition comprises detecting
3 that a difference in used capacity of a least full storage unit and the most full storage unit in
4 the plurality of storage units has exceeded a second threshold value, and the condition is
5 determined to be resolved when the difference is within the second threshold value.

1 47. The system of claim 46 wherein the program instructions when
2 executed by the processor, cause the processor to identify the most full storage unit as the
3 first storage unit.

1 48. The system of claim 43 wherein the program instructions when
2 executed by the processor, cause the processor to:
3 determine a storage unit from the plurality of storage units that is least full,
4 determine a storage unit from the plurality of storage units that is most full,
5 determine a difference in used capacity between the least full storage unit and
6 the most full storage unit, and

7 perform, the identification of a first storage unit from the plurality of storage
8 units from which data is to be moved, the identification of a file stored on the first storage
9 unit to be moved, the identification of a storage unit from the plurality of storage units for
10 storing the file, the move of the file from the first storage unit to the storage unit identified for
11 storing the file, and the repeating, only if the difference exceeds a pre-configured threshold
12 value.

1 49. The system of claim 43 wherein the program instructions when
2 executed by the processor, cause the processor to:
3 generate a score for each file included in a plurality of files stored on the first
4 storage unit, and
5 select a file, from the files stored on the first storage unit, a file with the
6 highest score as the file to be moved.

1 50. The system of claim 43 wherein the program instructions when
2 executed by the processor, cause the processor to:
3 generate a score for the storage units in the plurality of storage units, and
4 select a storage unit from the plurality of storage units with the highest score
5 as the storage unit for storing the file.

1 51. The system of claim 43 wherein the program instructions when
2 executed by the processor, cause the processor to:
3 determine a storage unit from the plurality of storage units that is least full,
4 determine a storage unit from the plurality of storage units that is most full,
5 determine a difference in used capacity between the least full storage unit and
6 the most full storage unit, and
7 repeat, the identification of a file stored on the first storage unit to be moved,
8 the identification of a storage unit from the plurality of storage units for storing the file, and
9 the move of the file from the first storage unit to the storage unit identified for storing the file,
10 only if the difference exceeds a pre-configured threshold value.

1 52. The system of claim 43 wherein the plurality of storage units
2 comprises at least one storage unit assigned to a first server and at least another storage unit
3 assigned to a second server distinct from the first server.

1 53. The system of claim 43 wherein an original file stored on the first
2 storage unit is not moved until all migrated files stored on the first storage unit have been
3 moved.

1 54. In a storage environment comprising a plurality of storage units
2 assigned to one or more servers, a system for performing capacity utilization balancing, the
3 system comprising:
4 at least one processor; and
5 a memory operatively coupled to the processor, the memory storing program
6 instructions that when executed by the processor, cause the processor to:
7 monitor a first group of storage units from the plurality of storage
8 units,
9 receive a first signal indicative of a condition,
10 determine, responsive to the first signal, a first storage unit from the
11 first group of storage units from which data is to be moved, and
12 move data from the first storage unit to one or more other storage units
13 in the first group of storage units until the condition is resolved.

1 55. The system of claim 54 wherein:
2 the first signal indicates that used storage capacity for a storage unit from the
3 first group of storage units has exceeded a first threshold; and
4 the program instructions when executed by the processor, cause the processor
5 to identify the storage unit whose used storage capacity has exceeded the first threshold as the
6 first storage unit.

1 56. The system of claim 54 wherein the program instructions when
2 executed by the processor, cause the processor to:
3 identify a file stored on the first storage unit to be moved,
4 identify a storage unit from the first group of storage units for storing the file,
5 move the file from the first storage unit to the storage unit identified for
6 storing the file, and
7 repeat, the identification of a file, identification of a storage unit, and the move
8 of the file, until the condition is determined to be resolved..

1 57. The system of claim 56 wherein the program instructions when
2 executed by the processor, cause the processor to:
3 generate a score for each file included in a plurality of files stored on the first
4 storage unit, and
5 select a file, from the files stored on the first storage unit, with the highest
6 score as the file to be moved.

1 58. The system of claim 56 wherein the program instructions when
2 executed by the processor, cause the processor to:
3 generate a score for the storage units in the first group of storage units, and
4 select a storage unit from the first group of storage units with the highest score
5 as the storage unit for storing the file.

1 59. The system of claim 54 wherein the program instructions when
2 executed by the processor, cause the processor to:
3 move a first file stored on the first storage unit to a first target storage unit
4 included in the first group of storage units, and
5 move a second file stored on the first storage unit to a second target storage
6 unit included in the first group of storage units, wherein the second target storage unit is
7 distinct from the first target storage unit.

1 60. The system of claim 54 wherein the program instructions when
2 executed by the processor, cause the processor to:
3 determine a storage unit from the first group of storage units that is least full,
4 determine a storage unit from the first group of storage units that is most full,
5 determine a difference in used capacity between the least full storage unit and
6 the most full storage unit, and
7 perform the determining the first storage unit step and the moving step only if
8 the difference exceeds a pre-configured threshold value.

1 61. The system of claim 54 wherein the program instructions when
2 executed by the processor, cause the processor to receive information indicative of storage
3 units from the plurality of storage units to be included in the first group of storage units.

1 62. The system of claim 54 wherein the first group of storage units
2 comprises at least one storage unit assigned to a first server and at least another storage unit
3 assigned to a second server distinct from the first server.

1 63. The system of claim 54 wherein original data stored on the first storage
2 unit is not moved until all migrated data stored on the first storage unit has been moved.

1 64. A system for balancing capacity utilization in a storage environment
2 comprising storage units, the system comprising::
3 means for detecting a condition indicating that capacity utilization balancing is
4 to be performed for a plurality of storage units;
5 means for identifying a first storage unit from the plurality of storage units
6 from which data is to be moved;
7 means for identifying a file stored on the first storage unit to be moved;
8 means for identifying a storage unit from the plurality of storage units for
9 storing the file;
10 means for moving the file from the first storage unit to the storage unit
11 identified for storing the file; and
12 means for repeating, identifying a file stored on the first storage unit to be
13 moved, identifying a storage unit from the plurality of storage units for storing the file, and
14 moving the file from the first storage unit to the storage unit identified for storing the file,
15 until the condition is determined to be resolved.

1 65. A system for performing capacity utilization balancing in a storage
2 environment comprising a plurality of storage units assigned to one or more servers, the
3 system comprising:
4 monitoring a first group of storage units from the plurality of storage units;
5 receiving a first signal indicative of a condition;
6 determining, responsive to the first signal, a first storage unit from the first
7 group of storage units from which data is to be moved; and
8 moving data from the first storage unit to one or more other storage units in
9 the first group of storage units until the condition is resolved.